Technical Paper Series
Macroeconomic Analysis and Tax Analysis Divisions
Congressional Budget Office
Washington, D.C.

# Can Permanent-Income Theory Explain Cross-Section Consumption Patterns?

John Sabelhaus
Jeff Groen

July 1997

1997-3

# CAN PERMANENT-INCOME THEORY EXPLAIN CROSS-SECTION CONSUMPTION PATTERNS?

John Sabelhaus
Jeff Groen*

April 1997

Abstract

The prediction that consumption-income ratios will decline as income rises in cross-section data is a feature of Friedman's [1957] Permanent Income Hypothesis and other intertemporal consumption-smoothing models. That prediction underlies our approach; we use longitudinal income data to predict consumption-income ratios across income groups, then compare those predictions to actual values. We show that models with a fixed propensity to consume out of permanent income cannot explain the skewness in annual consumption-income ratios. Allowing the long-run propensity to consume to decline with permanent income is the key to replicating actual consumption-income ratios.

JEL Codes: D12 Consumer Economics; Empirical Analysis. E21 Consumption; Saving.

*Tax Analysis Division, Congressional Budget Office, Washington, D.C. 20515. The views here do not necessarily reflect the opinion of the Congressional Budget Office. We presented an earlier version of this paper at the January, 1997 American Economic Association Meetings in New Orleans, La. We thank Frank Sammartino for useful comments and suggestions at every stage of this research. We also thank Eric Engen and seminar participants at the Congressional Budget

## I. Introduction

The prediction that consumption-income ratios will be negatively correlated with income in cross-section data is a well-known feature of Friedman's [1957] Permanent Income Hypothesis (PIH) and other intertemporal consumption-smoothing theories. Some families with low annual income have higher permanent income, so average consumption in low income groups will be greater than average income. At the top of the annual income distribution, the opposite holds: some families with high annual income have lower permanent income, so average consumption in high income groups will be less than average income.

The PIH and other consumption-smoothing models are important building blocks in economic theory, but the predicted relationship between annual consumption and permanent income cannot be tested directly because no data set exists that includes both family-level consumption and incomes over time. This paper therefore investigates the relationship between consumption and permanent income indirectly. We use longitudinal income data to predict the distribution of annual consumption-income ratios under various versions of the PIH and then compare those predictions to actual cross-section data. Thus we can discern which versions of the theory are able to reconcile the longitudinal income and cross-section expenditure data.

We show first that a simple version of the PIH in which the propensity to consume is constant across permanent income groups is able to explain only about half the skewness in cross-section consumption-income ratios. Basically, given a fixed propensity to consume out of permanent income, there is not enough variability in family-level incomes over time to explain why consumption-income ratios are so high for lower-income families, and so low for higher-income families. More realistic intertemporal models with uncertainty suggest that consumption is

determined by both permanent income and annual income, and therefore those models predict even less cross-section skewness than in the simple certainty-equivalence case.[1]

However, a reasonable extension of the simple case is able to replicate the observed spending patterns. If propensities to consume decrease with permanent income, the predicted skewness in annual consumption-income ratios increases, because families in the lowest permanent income groups (who are most likely to be in the lowest annual income groups) have higher consumption propensities. We use this logic in reverse and solve for the schedule of permanent consumption propensities that replicates the observed cross-section consumption-income ratios. As expected, the fitted permanent consumption propensities across permanent income groups are much less skewed than the annual consumption propensities, because (as in the simple case) some of the observed pattern of consumption-income ratios can be attributed to income variability.[2]

We then compare the estimated propensities to consume out of permanent income to other (implied) estimates from various wealth-based saving measures. Although we find much less skewness in permanent consumption propensities than in the annual estimates, there is still more skewness than would be consistent with measuring changes in wealth over time. The fitted propensity to consume for the bottom income decile is 127 percent, but wealth-change studies find values closer to 100 percent. For the top permanent income decile, the fitted propensity to consume is 73 percent, but the wealth-change studies find values of 85 percent or more.

---

[1] See, for example, Carroll [1997], Hubbard, Skinner, and Zeldes [1995], and Huggett and Ventura [1996]. For a comprehensive review of theory and facts about household-level saving and consumption behavior, see Browning and Lusardi [1996].

[2] The declining pattern of consumption propensities across permanent income groups is consistent with predictions of models that emphasize the redistribution built into real-world tax and transfer systems. See, for example, Hubbard, Skinner, and Zeldes [1995], Huggett and Ventura [1996], and Sabelhaus [1997].

Thus, the PIH and longitudinal income variability cannot completely explain the pattern of cross-section consumption-income ratios. One remaining (and likely) explanation is systematic income measurement error; people at the bottom of the income distribution may under report their income on surveys, thus overstating consumption-income ratios at the bottom. That proposition is supported by cross-section wealth holdings, because there is insufficient wealth (plus income) to account for reported consumption at low income levels. Therefore, either consumption is over reported, or income is under reported, and the latter seems more likely. If valid, that explanation also suggests problems with existing estimates of how income is distributed across the population, because the distribution of income reported in the cross-section expenditure survey is similar to the distribution in other (including income-oriented) surveys except at the highest income levels.

## II. Consumption-Income Ratios in Annual Survey Data

Although aggregate consumption is a large and relatively stable share of aggregate income over time, the PIH and other consumption-smoothing models predict that the ratio of consumption to annual income will vary systematically across income groups. This section documents the pattern of consumption-income ratios in the Consumer Expenditure Survey (CEX) for 1992, using various combinations of consumption concepts, distributional statistics, and age groupings.[3]

Several tabulations from the CEX are shown in Table 1. The distributional classifier is before-tax family income divided by the family size adjustment implicit in the Census poverty thresholds. The Census scale suggests, for example, that a family with two members requires only

[3] See the data appendix and Sabelhaus [1996] for a description of the CEX sample used in this study.

28 percent more income to reach the same level of well-being as a single individual, because of shared resources and economies of scale.[4] The first column is the ratio of total consumption to total (after-tax) income in each decile.[5] The pattern of consumption-income ratios is similar to other findings using the CEX (Sabelhaus [1993]; Feenberg, Mitrusi, and Poterba [1997]; Poterba [1989]). Families in the bottom decile spend 230 percent of their income, while families in the top decile spend only 64 percent. The pattern between across deciles is non-linear and convex—the bottom five deciles have negative saving rates, while the top five have positive rates.

The distributional pattern of consumption-income ratios across income groups is robust with respect to alternative measures of consumption. The measure in the second column excludes durable goods because those purchases are volatile and partly represent investment. For example, someone who earns $20,000 a year and buys a $10,000 car that will last five years is consuming $2,000 worth of car per year. If the durable purchase is included in total consumption, the ratio of consumption to income will be overstated. But even when the expenditure concept excludes durables, the skewness in consumption-income ratios across income groups persists. That is, the impact of durable purchases averages out across the many families within a given decile, some of whom have purchases and some of whom do not.

The third column reports median consumption-income ratios. If average consumption within a decile is strongly influenced by a few outliers, the median ratios are a better indicator of typical

---

[4] The Census scale adjustments for family sizes 2 through 9 are, respectively, 1.28, 1.57, 2.01, 2.38, 2.68, 3.04, 3.38, 4.04. All of the results here and throughout the paper are basically the same across three approaches we tested: no adjustment for family size, a per-adult adjustment, and the Census adjustment described and used in the paper.

[5] It is important to note that these are not average consumption-income ratios, they are ratios of average consumption to average income. The former can show even more skewness, particularly if very low-income families are included.

consumption behavior within the group. Though the median consumption-income ratio in the bottom decile is a bit lower than the mean ratio, the overall pattern remains. The typical family in the bottom decile spends 186 percent of disposable income, which is below the average of 230 percent. In the top decile the median and average ratios of consumption to income are identical at 64 percent.

The last three columns show that, after controlling for income, age has little explanatory effect. The overall consumption-income ratio is lower for middle-aged people (40 to 60) than it is for the young (<40) or old (60+). That observation seems consistent with life-cycle versions of consumption-smoothing theories—middle-aged people are at the point in their life when saving should be high. But, surprisingly, that finding disappears within any given income decile. For example, in the top decile, the young spend 65 percent of their income, the middle-aged 64 percent, and the old 61 percent. In general, the overall skewness in consumption-income ratios across income groups holds within any given age group, as well as in the aggregate. [6]

---

[6] These results should not be interpreted as evidence against the life-cycle model, for two reasons. First, we use the term "consumption" loosely—our cash expenditure concept includes items that are not consumption, such as investment in durable goods. Also, our measure omits important age-related consumption components such as the imputed rental value of housing and the value of employer- and government-provided medical services. A comprehensive measure of consumption that makes those adjustments (such as in Gokhale, Kotlikoff, and Sabelhaus [1996]) does exhibit a life-cycle pattern. Second, our measure of "income" is also limited. We do not include employer pension contributions or Social Security taxes in saving and we do include pension and Social Security benefits in income. If we compute consumption and income in a way consistent with life-cycle theory, the expected pattern would show up in the tables. The calculations here are focussed on reconciling only the cash-flow piece of saving.

## III. Can Permanent-Income Theory Explain the Skewness in Consumption?

The skewness of consumption-income ratios across annual income groups is a well-known characteristic of expenditure survey data. Does that skewness result mainly from consumption-smoothing behavior? The appropriate data set for addressing that question is a panel survey with annual consumption and annual income over a long period. Unfortunately, that data set does not exist. So we use panel data on incomes to measure income variability over time, then assess whether the pattern of variability in the income data together with various versions of consumption-smoothing behavior can explain the observed consumption-income skewness in the expenditure data.

The data used to measure income variability are from the Panel Survey of Income Dynamics (PSID). The sample covers the period from 1982 through 1991, and includes most of the 1991 sample members who were in the survey for the ten-year period.[7] Our income measure in the PSID is the same as in the CEX: total family income adjusted for family size. Changes in family size, holding total income constant, will change a family's relative income position.[8]

Our measure of "permanent" income is the average (adjusted) annual income over this ten-year period. We remove the effect of economy-wide real-income growth by indexing average incomes across the years. Thus, average income in each of the sample years is the same, but the

---

[7] We excluded people who left their parents' family and established new families during the period. See the data appendix for a more detailed description of the PSID sample used here.

[8] This implicitly assumes that adjusted family income is the right way to measure well-being across families. An alternative strategy is to restrict the sample to families that were intact over the ten years. That approach avoids the need to adjust income for family composition change. We prefer the approach in the paper because it uses much more of the sample. In any case, we used the restricted-sample approach in an earlier version of the paper and generated similar results.

income variability measures capture idiosyncratic movements (including progression through age-earnings profiles) for any given family.

Table 2 shows various statistics from the annual income and permanent income distributions in our PSID sample. The annual income distribution is based on all ten years of data—each family shows up ten times, based on their annual income in each year. (The same basic result can be derived using any year's annual data, but the decile breaks and averages within deciles are slightly sensitive to the exact year chosen, especially within the thin upper tail of the distribution). The main message of Table 2 is that classifying people by annual income does indeed produce a more dispersed distribution than classifying people by permanent income—annual decile breaks and average incomes within deciles are lower at the bottom and higher at the top in the annual distributions.

Table 3 shows the effects of cross-tabulating observations by permanent and annual income deciles. Again, each observation shows up in Table 3 ten times to avoid thin distributions in any given year. The distribution of annual incomes within any permanent income decile (column) can be read by moving down the annual deciles (rows). For example, a family whose permanent income places it in the bottom decile (less than $7,600, Table 2) has a 69.6 percent chance of being in the bottom annual decile (less than $6,420). It has a 23.8 percent chance of being in the second annual decile (between $7,660 and $11,750), a 4.2 percent chance of being in the third annual decile, and a much smaller chance of being in any of the fourth through tenth annual deciles.

A similar decomposition for annual income groups can be read off the rows of Table 3. Again, a family whose annual income places them in the bottom decile in a given year has a 69.6 percent chance of being in the bottom permanent decile, a 17.1 percent chance of being in the second, 6.7 percent in the third, and so on. Putting the numbers in context, average consumption in

the bottom annual decile in any given year will equal (0.696)*(average consumption in the bottom permanent income group) plus (0.171)*(average consumption in the second permanent income group) plus (0.067)*(average consumption in the third permanent income group), and so on.

Table 3 suggests that families' decile rankings are relatively stable, particularly among the very poor and very rich. About 70 percent of the permanent poor are annual poor, and about 70 percent of the permanent rich are annual rich. Almost all income variability is restricted to plus or minus one decile. There is virtually no overlap between the extremes of the permanent and annual income distributions, though it is much more likely for a permanent-rich person to have a bad year and show up in the lower annual deciles than for a permanent-poor person to have a good year and show up in the higher annual deciles.[9]

In our first experiment with these income variability estimates, we simply compute the pattern of annual consumption-income ratios, assuming that the simple PIH holds. Consumption in a given annual income decile is the sum of consumption over all families in that annual income decile, assuming that each family's consumption is proportional to its permanent income. For annual income group k, total consumption is the sum of consumption across families whose annual income is in the appropriate range (all $i \in k$), that is,

$$(1) \qquad C_k = \sum_{i \in k} \beta \, y_i \, ,$$

[9] The impression from Table 3 is quite different from a similar table (4-8) in Fullerton and Rogers [1993]. Their table cross-classifies people by annual and lifetime incomes, and shows significantly more dispersion. There are a few reasons why the tables differ. Our table uses data from single families over ten years, whereas the Fullerton and Rogers table uses data from estimated age-income profiles across eleven discrete groups over entire lifetimes. Also, their age-income profiles include the effect of economy-wide real wage growth, which we eliminate by construction.